

U.S. Non-Provisional Patent Application

Attorney Docket No.: 200308565-1

Title:

MEMORY QUALITY ASSURANCE

Inventors:

Ken Gary Pomaranski

406 Las Cruces, Ct.

Roseville, Ca., 95747

Citizenship: United States

Andy Harvey Barr

117 Montgomery St.

Roseville, Ca., 95678

Citizenship: United States

Dale John Shidla

407 Las Cruces, Ct.

Roseville, Ca., 95747

Citizenship: United States

MEMORY QUALITY ASSURANCE

BACKGROUND

[0001] Computer systems, and other electronic systems may have electronic memory.

Some of this may be referred to as “main memory”. Main memory may be built from, for example, dynamic random access memory (DRAM) chips. The DRAMs may be organized, for example, onto memory boards or partitioned into dual in line memory modules (DIMMS). Memory may experience errors like transient single bit errors, multi-bit errors, stuck-at single bit errors, and the like, which can negatively impact the systems in which the memory is located.

[0002] An operating system or other control system associated with the computer or electronic system may regard the memory as a logical pool of available memory. The operating system may virtualize the available memory so that it can be managed, shared, accessed and so on by various operating system instances (e.g., applications, threads, processes, programs). Thus, physical memory addresses may be translated to virtual memory addresses and vice versa by one or more logics.

[0003] Memory usage may vary during system operation depending, for example, on the type, number, size and so on, of applications running on a system. The variance may lead to some memory areas being used frequently while others are used less frequently. If a system is configured to detect memory errors, then errors in more frequently used areas may be more likely to be discovered, accounted for, handled, and so on than errors in less frequently used areas. Thus, errors in less frequently used areas may go undiscovered and may eventually evolve into catastrophic errors as the errors accumulate.

[0004] Conventional systems may have employed application level “software memory scrubbing” techniques in an attempt to exercise memory, and to discover, account for, and perhaps correct certain memory errors. However, conventional software memory scrubbing may negatively impact system performance by disturbing (e.g., interrupting, halting, messaging) an operating system, control system, or user level application and/or by consuming non-memory resources (e.g., processor cycles, file table entries, process table entries) that would otherwise be available for operating system instances. Furthermore, conventional software memory scrubbing may not be able to access all or even substantially all of the memory in a system if some memory is locked by an operating system, control

system, operating system instance, or the like. Additionally, some memory may not be accessible if memory has been partitioned by, for example, an operating system. Thus, errors may still go undetected and may accumulate in areas that software memory scrubbing does not reach. Even if a conventional software scrubber detects a memory area that may have

5 suspect qualities (e.g., parity error detected in location), the application may be limited in its response to the detected error. For example, a software memory scrubbing application may log an error location. The log may then be read by a separate diagnostic software application after a system shutdown and reboot.

10 BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate various example systems, methods, and so on that illustrate various example embodiments of aspects of the invention. It will be appreciated that the illustrated element boundaries (e.g., boxes, groups of boxes, or other shapes) in the figures represent one example of the boundaries. One of ordinary skill in the art will appreciate that one element may be designed as multiple elements or that multiple elements may be designed as one element. An element shown as an internal component of another element may be implemented as an external component and vice versa. Furthermore, elements may not be drawn to scale.

20 [0006] Figure 1 illustrates an example memory quality assurance system.

[0007] Figure 2 illustrates an example memory quality assurance logic.

[0008] Figure 3 illustrates another example memory quality assurance logic.

[0009] Figure 4 illustrates an example memory quality assurance method.

[0010] Figure 5 illustrates another example memory quality assurance method.

25 [0011] Figure 6 illustrates an example computing environment in which example memory quality assurance systems and methods can operate.

[0012] Figure 7 illustrates an example image forming device in which example memory quality assurance systems and methods can operate.

30 [0013] Figure 8 illustrates an example operating system transparent system for on-the-fly memory testing.

- [0014] Figure 9 illustrates an example memory quality assurance method.
- [0015] Figure 10 illustrates another example memory quality assurance method.
- [0016] Figure 11 illustrates an example operating system transparent method for on-the-fly memory testing.

5

DETAILED DESCRIPTION

[0017] This application describes example systems, methods, computer-readable mediums and so on associated with assuring the quality of electronic memory without disturbing operating system instances. The example systems, methods, computer-readable 10 mediums and so on facilitate exercising memory to detect memory errors, where the exercising occurs in parallel and/or substantially in parallel with normal system operation without disturbing normal system operation (e.g., halting an application whose memory is being tested). In one example, the exercising does not engage user applications, operating systems or other similar control systems and thus does not interfere with the performance of 15 such applications, operating systems and so on. In one example, memory errors can be detected, predicted, and/or accounted for proactively without involving a user level application or operating system.

[0018] The following includes definitions of selected terms employed herein. The definitions include various examples and/or forms of components that fall within the scope of 20 a term and that may be used for implementation. The examples are not intended to be limiting. Both singular and plural forms of terms may be within the definitions.

[0019] “Computer-readable medium”, as used herein, refers to a medium that participates 25 in directly or indirectly providing signals, instructions and/or data. A computer-readable medium may take forms, including, but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media may include, for example, optical or magnetic disks and so on. Volatile media may include, for example, optical or magnetic disks, dynamic memory and the like. Transmission media may include coaxial cables, copper wire, fiber optic cables, and the like. Transmission media can also take the form of electromagnetic radiation, like those generated during radio-wave and infra-red data communications, or take 30 the form of one or more groups of signals. Common forms of a computer-readable medium include, but are not limited to, an application specific integrated circuit (ASIC), a compact

disc (CD), a digital video disk (DVD), a random access memory (RAM), a read only memory (ROM), a programmable read only memory (PROM), an electronically erasable programmable read only memory (EEPROM), a disk, a carrier wave, a memory stick, a floppy disk, a flexible disk, a hard disk, a magnetic tape, other magnetic media, a CD-ROM,

5 other optical media, punch cards, paper tape, other physical media with patterns of holes, an EPROM, a FLASH-EPROM, or other memory chip or card, and other media from which a computer, a processor or other electronic device can read. Signals used to propagate instructions or other software over a network, like the Internet, can be considered a “computer-readable medium.”

10 [0020] “Logic”, as used herein, includes but is not limited to hardware, firmware, software and/or combinations of each to perform a function(s) or an action(s), and/or to cause a function or action from another component. For example, based on a desired application or needs, logic may include a software controlled microprocessor, discrete logic like an ASIC, a programmed logic device, a memory device containing instructions, or the like. Logic may
15 also be fully embodied as software. Where multiple logical logics are described, it may be possible to incorporate the multiple logical logics into one physical logic. Similarly, where a single logical logic is described, it may be possible to distribute that single logical logic between multiple physical logics.

20 [0021] “Signal”, as used herein, includes but is not limited to one or more electrical or optical signals, analog or digital, one or more computer or processor instructions, messages, a bit or bit stream, or other means that can be received, transmitted and/or detected.

[0022] “Software”, as used herein, includes but is not limited to, one or more computer or processor instructions that can be read, interpreted, compiled, and/or executed and that cause a computer, processor, or other electronic device to perform functions, actions and/or behave
25 in a desired manner. The instructions may be embodied in various forms like routines, algorithms, modules, methods, threads, and/or programs including separate applications or code from dynamically linked libraries. Software may also be implemented in a variety of executable and/or loadable forms including, but not limited to, a stand-alone program, a function call (local and/or remote), a servelet, an applet, instructions stored in a memory, part
30 of an operating system or other types of executable instructions. It will be appreciated by one of ordinary skill in the art that the form of software may depend on, for example, requirements of a desired application, the environment in which it runs, and/or the desires of

a designer/programmer or the like. It will also be appreciated that computer-readable and/or executable instructions can be located in one logic and/or distributed between two or more communicating, co-operating, and/or parallel processing logics and thus can be loaded and/or executed in serial, parallel, massively parallel and other manners.

5 [0023] Suitable software for implementing the various components of the example systems and methods described herein include programming languages and tools like Java, Pascal, C#, C++, C, CGI, Perl, SQL, APIs, SDKs, assembly, machine, firmware, microcode, and/or other languages and tools. Software, whether an entire system or a component of a system, may be embodied as an article of manufacture and maintained as part of a computer-
10 readable medium as defined previously. Another form of the software may include signals that transmit program code of the software to a recipient over a network or other communication medium.

[0024] “Data store”, as used herein, refers to a physical and/or logical entity that can store data. A data store may be, for example, a database, a table, a file, a list, a queue, a heap, a
15 memory, a register, and so on. A data store may reside in one logical and/or physical entity and/or may be distributed between two or more logical and/or physical entities.

[0025] An “operable connection”, or a connection by which entities are “operably connected”, is one in which signals, physical communication flow, and/or logical communication flow may be sent and/or received. Typically, an operable connection
20 includes a physical interface, an electrical interface, and/or a data interface, but it is to be noted that an operable connection may include differing combinations of these or other types of connections sufficient to allow operable control.

[0026] Some portions of the detailed descriptions that follow are presented in terms of algorithms and symbolic representations of operations on data bits within a memory. These
25 algorithmic descriptions and representations are the means used by those skilled in the art to convey the substance of their work to others. An algorithm is here, and generally, conceived to be a sequence of operations that produce a result. The operations may include physical manipulations of physical quantities. Usually, though not necessarily, the physical quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined,
30 compared, and otherwise manipulated in a logic and the like.

[0027] It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be borne in mind, however, that these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these

5 quantities. Unless specifically stated otherwise, it is appreciated that throughout the description, terms like processing, computing, calculating, determining, displaying, or the like, refer to actions and processes of a computer system, logic, processor, or similar electronic device that manipulates and transforms data represented as physical (electronic) quantities.

10 [0028] **Figure 1** illustrates a memory quality assurance system **100**. The system **100** facilitates testing memory substantially in parallel with normal system operation without operating system or user application level involvement. The system **100** selectively and temporarily mirrors then logically replaces main memory locations with spare memory locations while maintaining normal system operation. The system **100** then tests or has tested 15 on its behalf the logically replaced main memory location(s). While the testing is in progress, memory accesses intended for those memory location(s) being tested are redirected to spare memory location(s). In some cases, where the memory testing reveals an error, the main memory location(s) may be logically removed from the main memory and the temporary logical replacement(s) extended to logically replace the tested memory location(s).

20 [0029] The system **100** accesses main memory. The main memory may be allocated, for example, into various sets of memory associated with various operating system instances (e.g., programs, applications, threads, processes). Thus, a first set of memory **10** may include, for example, memory locations **12**, and **14** through **16**. The first set of memory **10** may be, for example, a relatively small set of memory (e.g., 10K) or a relatively larger set of 25 memory (e.g., 16G). Similarly, a second set of memory **20** may include, for example, memory locations **22**, and **24** through **26** while a third set of memory **30** may include, for example, memory locations **32**, and **34** through **36**. While three sets of memory are illustrated it is to be appreciated that at various points in time that various computer systems may have a greater and/or lesser number of sets of memory of various sizes allocated to 30 various threads, processes, operating system instances, and so on.

[0030] In one example, the system **100** may also access a separate set of memory, which may be referred to as “spare memory” **40**. The spare memory **40** may include, for example,

memory locations **42**, and **44** through **46**. Once again, it is to be appreciated that the spare memory **40** may be of various sizes. While the spare memory **40** is illustrated separate from the main memory, it may physically be a part of the main memory while being logically separated out. In one example, the spare memory **40** is memory that is known to have quality attributes exceeding a pre-determined, configurable threshold (e.g. no memory errors in last 1 million accesses).

[0031] The system **100** may include a memory mapping logic **110** that can interact with a memory quality assurance logic **120** to facilitate memory quality assurance testing. The memory mapping logic **110** can be configured to provide access to memory locations. By

5 way of illustration, a processor **130** that is processing a set of applications (e.g., A₁ **140**

through A_n **144**, n being an integer) may seek access to various memory locations. For

example, the processor **130** may wish to perform a memory accessing operation like an

input/output operation (i/o) to a memory location. Performing the i/o may include the

processor **130** sending a memory address to the memory mapping logic **110**. The memory

15 mapping logic **110** may then resolve the address and send one or more signals to a physical

memory location to perform the i/o. In one example, the memory mapping logic **110** can be

configured to redirect a memory accessing operation intended for a first memory location to a

second memory location. Thus, an i/o intended for a main memory location can be directed

to a spare memory location. In one example, the spare memory **40** may be located in the

20 memory mapping logic **110** while in another example the spare memory **40** may be located in the
memory quality assurance logic **120**.

[0032] By way of further illustration, application A₁ **140** may want to write a value to

memory. The processor **130** may therefore perform an output operation. The output

operation can include sending a memory address to the memory mapping logic **110**. The

25 memory mapping logic **110** may resolve the memory address and complete the output

operation to memory location **12** in memory set **10**, which may be associated with application

A₁ **140**. Similarly, application A₂ **142** may want to read a value from memory. The

processor **130** may therefore perform an input operation. The input operation can include

sending a memory address to the memory mapping logic **110**, which resolves the memory

30 address and completes the input operation from memory location **26** in memory set **20**, where

memory set **20** is associated with application A₂ **142**. However, if the memory mapping

logic **110** has been reconfigured by the memory quality assurance logic **120**, then the input

operation may have taken a value from memory location **42** in spare memory **40** rather than from memory location **26** in memory set **20**. Thus, a first memory location (e.g. **26**) can be logically replaced by a second memory location (e.g. **42**) leaving the first memory location logically isolated and available for testing.

5 [0033] The memory quality assurance logic **120** may track memory ownership relationships between operating system instances (e.g., applications) and physical memory locations. Additionally the memory quality assurance logic **120** may store memory redirection data like main memory addresses that have been logically replaced by spare memory addresses and the relationships between them. Thus, in one example, the memory
10 quality assurance logic **120** may include one or more data stores configured to store one or more, of a memory freshness data, a memory quality data, an operating system instance to physical memory location relationship data, and a memory reconfiguration data. In another example, the memory quality assurance logic **120** may include a microprocessor, a memory and a non-volatile memory. The non-volatile memory may store, for example, memory
15 location freshness data (e.g., how recently it has been accessed and/or error checked) and/or memory location quality data (e.g., error rate, error types). In another example, the memory quality assurance logic **120** may be operably connected to one or more data stores configured to store one or more, of a memory freshness data, a memory quality data, an operating system instance to physical memory location relationship data, and a memory reconfiguration data.
20 [0034] The memory quality assurance logic **120** may be configured to select a memory location to error check. The selection may be made by methods including, but not limited to, linearly (e.g., memory locations chosen in order), round-robin (e.g., memory locations chosen in order, then loop back to first location after last location is chosen and continue), randomly, most frequently used, least frequently used, and so on. The memory quality assurance logic
25 **120** may also be configured to select a location in spare memory **40** to participate in the error checking. The memory quality assurance logic **120** may be configured to send one or more signals that cause the copying of the contents of the memory location to be error checked to the selected spare memory location. The memory quality assurance logic **120** may also send one or more signals to the memory mapping logic **110** that cause future memory accessing
30 requests initially destined for the memory location to be error checked to be routed to the selected spare memory location. The memory quality assurance logic **120** may also send one or more signals that initiate error checking the memory location to be error checked. For

example, the memory quality assurance logic **120** may send a diagnostic initiating signal to a memory board or memory chip associated with the memory location to be error checked.

The memory board or memory chip may then perform diagnostics (e.g., error checking) on the memory location. The diagnostics may be stored, for example, in hardware, firmware,

5 and/or software on the memory board or chip. These diagnostics may facilitate determining whether a memory location is experiencing memory errors (e.g., parity errors, stuck bit errors). Note that performing these diagnostics does not engage the operating system or user level applications. The operating system, control system, user applications, operating system instances and so on that access the memory may not even be aware that the diagnostics are
10 being performed. Thus, the error checking may be done on-the-fly (e.g., during normal system operation without a halt or reboot) while remaining transparent (e.g., not halting or consuming resources) to operating system instances.

[0035] The results of the diagnostics can be reported, for example, to the memory quality assurance logic **120**. Based on the results of the diagnostics, the memory quality assurance

15 logic **120** may determine that the memory location has quality attributes exceeding a pre-determined configurable quality threshold (e.g., passed error checking) and thus send one or more signals to the memory mapping logic **110** to logically return the tested memory location to main memory and/or to reestablish a relationship between the memory location and an application, for example. Thus, subsequent memory access requests initially destined for the
20 memory location that was error checked will be delivered to the memory location that was error checked rather than to the spare memory location. Similarly, based on the results of the diagnostics, the memory quality assurance logic **120** may determine that the memory location has quality attributes falling below a pre-determined, configurable quality threshold (e.g., failed error checking). Thus the memory quality assurance logic **120** may decide to logically
25 remove the memory location from main memory and send zero or more signals to the memory mapping logic **110** so that subsequent memory accessing requests initially destined for the memory location that was error checked will continue to be delivered to the spare memory location. In another example, the memory quality assurance logic **120** may be configured to identify an alternate memory location in the memory set in which the memory
30 location to be tested is located and send one or more signals to the memory mapping logic **110** so that subsequent memory accessing requests initially destined for the memory location

that was error checked are delivered to the alternate memory location. In this way, the pool of spare memory may be preserved for future testing purposes.

[0036] **Figure 2** illustrates a system **200** that includes an example memory quality assurance logic **210**. The memory quality assurance logic **210** may be configured to interact with a memory mapping logic **220** to facilitate actions including, but not limited to, examining memory, exercising memory, detecting memory errors, and handling memory address remapping to mitigate the effects of memory errors. In one example, the memory mapping logic **220** may include a crossbar that facilitates resolving and/or redirecting memory addresses. In another example, the memory mapping logic **220** may include one or more programmable address translation tables that facilitate resolving and/or redirecting memory addresses. Thus, in one example, the memory quality assurance logic **210** may reconfigure the memory mapping logic **220** by reprogramming the crossbar. In another example, the memory quality assurance logic **210** may reconfigure the memory mapping logic **220** by reprogramming one or more entries in one or more address translation tables. While a crossbar and an address translation table are described, it is to be appreciated that other address mapping and/or resolving apparatus, methods, and data stores may be employed.

[0037] The memory quality assurance logic **210** may store information about a physical memory space (e.g., a main memory space). For example, the main memory space may include memory locations M1 **232**, M2 **234**, and M3 **236** through M_x **238**, x being an integer. The main memory space may be relatively small (e.g., 1K in an embedded system) or relatively large (e.g., 64 TB in a server). The memory quality assurance logic **210** may also store information about a “spare memory” space. In one example, the spare memory space may be located in a separate set of memory chips, boards and so on, while in another example the spare memory space may be logically partitioned from memory chips, boards and so on associated with the main memory space. Thus, the memory quality assurance logic **210** may include and/or be operably connected to one or more data stores configured to store one or more of, a memory freshness data, a memory quality data, an operating system instance to physical memory location relationship data, and a memory reconfiguration data.

[0038] The memory quality assurance logic **210** can be configured to control (e.g., program) the memory mapping logic **220** so that memory access requests intended for a main memory location can be redirected to a spare memory location. Similarly, the memory

quality assurance logic **210** can be configured to copy, or have the memory mapping logic **220** or some other component copy, or to control copying the contents of a main memory location between a spare memory location. Similarly, the memory quality assurance logic **210** can be configured to copy, or to cause another component to copy, or to control copying the contents of a first main memory location between a second main memory location.

[0039] Thus, in one example, the memory quality assurance logic **210** may be configured to select a memory location to test. The memory location may be chosen by methods like linear, round-robin, random, most frequently used, least frequently used, and so on.

Similarly, the memory quality assurance logic **210** may be configured to select a spare

10 memory location to hold the contents of the memory location to be tested. Once again, the spare memory location can be chosen by various methods (e.g., linear, random, round-robin).

[0040] The memory quality assurance logic **210** can also be configured to copy, control, and/or to initiate copying the contents of the memory location to be tested to the spare memory location. For example, the memory quality assurance logic **210** may send a signal to 15 the memory mapping logic **220** that causes a direct memory transfer between the two memory locations.

[0041] The memory quality assurance logic **210** can also be configured to selectively reconfigure (e.g., reprogram) the memory mapping logic **220** so that a memory accessing request (e.g., an i/o request) to a main memory location (e.g., M1 **232**) that has been selected 20 to test will be redirected to the selected spare memory location (e.g., S1 **242**).

[0042] The memory quality assurance logic **210** can also be configured to run and/or initiate the running of tests (e.g., functional, electrical) on the memory location to be tested.

The memory quality assurance logic **210** can also be configured to perform actions like storing, analyzing, and reporting the results of the testing. The memory quality assurance

25 logic **210** can also be configured to selectively respond to the memory test results. For example, if the tests reveal that a memory location should be logically removed from the active main memory pool, then the memory quality assurance logic **210** may make the temporary remapping in the memory mapping logic **220** more permanent and/or may establish a more permanent remapping to another memory location. This effectively 30 logically removes the tested main memory location from main memory and replaces it with a different memory location. In one example, the logical removal and replacement can be accomplished without interacting with an operating system, user application, or so on.

Additionally, a running operating system, user application, and the like may not even be aware that the testing, removal and replacement occurred.

[0043] Thus, examining **Figure 2**, consider a memory test that proceeds linearly from M1 232 through Mx 238. At a first point in time, the contents of M1 232 may be copied to S1 242. Then, the memory quality assurance logic 210 may reprogram the memory mapping logic 210 to redirect memory accessing requests for M1 232 to S1 242. Then, the memory quality assurance logic 210 may perform and/or initiate memory tests on memory location M1 232. After testing M1 232, if the memory location exhibits quality attributes that exceed a pre-determined, configurable quality threshold, then the memory quality assurance logic 210 may copy or initiate the copying of the contents of S1 242 back to M1 232 and reprogram the memory mapping logic 210 so that memory accessing requests for M1 232 are no longer delivered to S1 242 but are subsequently delivered to M1 232. The memory quality assurance logic 210 may be configured to step through the main memory pool so that all and/or substantially all of the main memory pool is eventually tested. For example, after testing M1 232, the memory quality assurance logic 210 may mirror, swap, and test M3 236, then Mx 238, then M2 234, and so on.

[0044] Suppose that during the testing the memory quality assurance logic 210 determines that the quality of M2 234 has fallen below a pre-determined, configurable quality threshold. The memory quality assurance logic 210 may then cause a spare memory location 20 to more permanently take the place of M2 234. This is the situation depicted in **Figure 3** where a replacement memory location has logically replaced a logically removed main memory location. “Logically replaced” means that memory accessing requests for the tested memory location will access the replacing memory location. “Logically replaced” does not mean that physical memory apparatus is physically moved from one place to another. Similarly, “logically removed” means that a memory location is no longer accessed by operating system and/or user level applications that are not aware of the logical replacement. The physical apparatus associated with the memory location need not be physically removed from the main memory to effect the logical removal.

[0045] **Figure 3** illustrates a system 300 that includes an example memory quality assurance logic 310 that has interacted with a memory mapping logic 320 to logically replace a main memory location M2 334 with a spare memory location S1 342. The memory quality assurance logic 310 may have stepped through the main memory (e.g., locations M1 332, M2

334, and M3 336 through M_x 338) and determined that the quality of M2 334 had fallen below a pre-determined, configurable quality threshold. While performing the memory tests, the memory quality assurance logic 310 may have been using spare memory location S1 342 for logically replacing (e.g., mirroring and swapping) the contents of main memory locations.

5 Thus, when M2 334 was determined to be in condition for replacement, the memory quality assurance logic 310 may have reprogrammed the memory mapping logic 320 so that future memory accessing requests for M2 334 were delivered to S1 342 and so that M2 334 was logically removed from the main memory pool. Therefore, for subsequent memory quality assurance testing, the memory quality assurance logic 310 may select between the remaining 10 spare memory locations (e.g., S2 344, and S3 346 through S_y 348, y being an integer).

[0046] Example methods may be better appreciated with reference to the flow diagrams of **Figures 4 and 5**. While for purposes of simplicity of explanation, the illustrated methodologies are shown and described as a series of blocks, it is to be appreciated that the methodologies are not limited by the order of the blocks, as some blocks can occur in 15 different orders and/or concurrently with other blocks from that shown and described. Moreover, less than all the illustrated blocks may be required to implement an example methodology. Furthermore, additional and/or alternative methodologies can employ additional, not illustrated blocks. In one example, methodologies are implemented as processor executable instructions and/or operations stored on a computer-readable medium.

20 [0047] In the flow diagrams, blocks denote “processing blocks” that may be implemented, for example, in software. Additionally and/or alternatively, the processing blocks may represent functions and/or actions performed by functionally equivalent circuits like a digital signal processor (DSP), an ASIC, and the like.

[0048] A flow diagram does not depict syntax for any particular programming language, 25 methodology, or style (e.g., procedural, object-oriented). Rather, a flow diagram illustrates functional information one skilled in the art may employ to fabricate circuits, generate software, or use a combination of hardware and software to perform the illustrated processing. It will be appreciated that in some examples, program elements like temporary variables, routine loops, and so on are not shown. It will be further appreciated that 30 electronic and software applications may involve dynamic and flexible processes so that the illustrated blocks can be performed in other sequences that are different from those shown and/or that blocks may be combined or separated into multiple components. It will be

appreciated that the processes may be implemented using various programming approaches like machine language, procedural, object oriented and/or artificial intelligence techniques.

[0049] **Figure 4** illustrates an example memory quality assurance method **400**. The method **400** may include, at **410**, selecting and/or identifying a memory location to test.

5 While method **400** illustrates selecting and/or identifying a memory location to test, it is to be appreciated that in one example the address of the memory location to test may be provided to the method **400**. The memory location may be, for example, a main memory location. The main memory location may be located, for example, in a DRAM. The DRAM may be located on a DIMM. The board on which the DRAM/DIMM is located or plugged into may 10 have built in diagnostics that are stored, for example, in firmware, hardware, and/or software. The location to test may be selected by methods including, but not limited to, linear, random, most frequently used, least frequently used, most recently exhibiting an error, least recently exhibiting an error, and so on.

[0050] The method **400** may also include, at **420**, logically replacing (e.g., mirroring and 15 swapping) the memory location to test with a spare memory location. The mirroring may involve, for example, copying the contents of the memory location to test to the spare memory location. At **420**, the method **400** may also include redirecting memory access requests (e.g., i/o requests) from the memory location to test to the spare memory location. The redirecting may be performed by, for example, reprogramming a crossbar and/or a 20 memory address translation table. Thus, at **430**, the memory location to test is logically isolated and can be tested while memory access requests are diverted to the spare memory location. The tests may include, for example, electrical, functional, parity, marching one, marching zeroes, stripe, “worst-case” pattern and other tests.

[0051] At **440**, a determination is made concerning whether to logically remove the 25 memory location that is tested at **430**. If the determination at **440** is Yes, that the memory location should be logically removed from main memory, then at **450** the logical removal may be completed. In one example, a temporary diversion from the memory location to the spare memory location may be made more permanent. In another example, the contents of the spare memory location may be mirrored to a second main memory location and an 30 additional address remapping performed so that memory accessing requests intended for the tested and failed main memory location are directed to the second memory location. If the determination at **440** is No, that the memory location should not be logically removed, then

the contents of the spare memory location can be copied back into the main memory location and the redirecting undone so that subsequent memory accessing requests are delivered to the main memory location and not the spare memory location.

[0052] While **Figure 4** illustrates various actions occurring in serial, it is to be
5 appreciated that various actions illustrated in **Figure 4** could occur substantially in parallel.
By way of illustration, a first process could select units to test and/or prepare a set of units to
test so that the next unit to test is available substantially immediately. The set of units to test
may be stored in a data store. Similarly, a second process could mirror and swap (e.g.,
logically replace) locations to test with spare memory locations. If multiple locations to test
10 have been identified by the first process, and if multiple spare memory locations are
available, then multiple tests may be run by one or more third processes substantially in
parallel. While three processes are described, it is to be appreciated that a greater and/or
lesser number of processes could be employed and that lightweight processes, regular
processes, threads, and other approaches could be employed.

15 [0053] **Figure 5** illustrates an example memory quality assurance method **500**. The
method **500** includes a two step process for memory quality assurance testing where a first
test may be employed to identify a suspect location and a second test may be employed to
more rigorously test a suspect location. Method **500** includes, at **510**, selecting a memory
location to test. The memory location may be selected by methods like, linear, round-robin,
20 random, most frequently used, least frequently used, and so on. While method **500** illustrates
selecting and/or identifying a memory location to test, it is to be appreciated that in one
example the address of the memory location to test may be provided to the method **500**. At
520, the contents of the memory location are mirrored into a spare memory location that is
known to have a quality level above a pre-determined configurable threshold. Also at **520**,
25 memory addressing is reprogrammed so that memory accesses intended for the selected
memory location are directed to the mirroring memory location. At **530**, a method for testing
the memory location is selected. Test methods may include, but are not limited to, parity
testing, stripe testing, marching one testing, marching zeroes testing, “worst case” pattern
testing, and so on. The test method may be selected based on factors like a previous quality
30 level for the memory location, how frequently the memory location has been used, the
criticality of an application to which the memory location has been allocated, the time
available to test the location, and so on. Once again, while method **500** illustrates selecting a

test method at **530** and selecting a secondary test method at **552**, it is to be appreciated that the test methods could be provided to method **500**.

[0054] At **540**, the memory location is tested. While a single memory location is described, it is to be appreciated that at **510** that two or more memory locations could be selected, that at **520** the two or more memory locations could be mirrored and swapped, that at **530** two or more test methods could be selected and that at **540** two or more memory locations could be tested, in serial, in parallel, and/or substantially in parallel. The parallel testing can be facilitated, in one example, by selecting at **530** a test method that is stored on a device associated with the memory location. For example, a memory location may be located in a DRAM on a memory board that has built in memory testing routines that can test multiple locations. Thus, two or more memory locations on the memory board may be tested serially, in parallel, and/or substantially in parallel. While the method **500** illustrates testing the location at **540**, it is to be appreciated that in one example the method **500** may initiate testing at **540** and be provided with test results concerning the tested memory location.

[0055] At **550**, a determination is made concerning whether a memory location is suspect. That is, did the memory testing routine(s) report that the memory location exhibited quality attributes that fell below a pre-determined configurable threshold. The threshold may be pre-determined and configurable to facilitate various levels of testing. By way of illustration, at a first time a system may be under a first relatively lighter load that makes more spare memory available for testing. Thus a first higher level of testing with a higher degree of parallelism may be undertaken. At a second time the system may be under a second relatively heavier load that makes less spare memory available for testing. Thus a second lower level of testing with a lower degree of parallelism may be undertaken. If the determination at **550** is No, then processing continues at **560**. But if the determination at **550** is Yes, then at **552** a secondary test method may be selected. The secondary test method may be selected to exercise the memory location in a manner that may uncover errors associated with the suspect attributes. The secondary method may then be employed at **554**.

[0056] At **560**, a determination is made concerning whether to logically remove the main memory location from the pool of available memory. Results from the test at **540** and/or the test at **554** may be considered when making the determination. If the determination at **560** is Yes, then at **570** the memory location is logically removed. Otherwise, if the determination at **560** is No, then at **580**, the memory location is logically returned to the pool of available

memory and/or re-associated with an operating system instance with which it had been previously associated. The logical removal and/or restoration may be effected by, for example, reprogramming a crossbar, reprogramming an address translation table, and so on.

[0057] While **Figure 5** illustrates various actions occurring in serial, it is to be appreciated that various actions illustrated in **Figure 5** could occur substantially in parallel. By way of illustration, a first process could select memory locations to test, a second process could select test methods for the memory locations, a third process could initiate and monitor the testing of memory locations, a fourth process could determine whether secondary testing is desired, a fifth process could select secondary test methods, a sixth process could initiate and monitor the secondary testing and a seventh process could determine whether to logically remove tested memory locations. While seven processes are described, it is to be appreciated that a greater and/or lesser number of processes could be employed and that lightweight processes, regular processes, threads, and other approaches could be employed.

[0058] In one example, a computer-readable medium may store processor executable instructions operable to perform a method that includes selecting a first memory location to test from a first set of memory. The method may also include selecting a second memory location to logically replace the first memory location during testing and copying the contents of the first memory location to the second memory location. The method may also include logically replacing the first memory location with the second memory location by reconfiguring address resolving means. The method may also include initiating testing of the first memory location and selectively logically replacing the first memory location with the second memory location based, at least in part, on the results of the testing. While one method is described, it is to be appreciated that other computer-readable mediums could store other example methods described herein.

[0059] **Figure 6** illustrates a computer **600** that includes a processor **602**, a memory **604**, and input/output ports **610** operably connected by a bus **608**. Executable components of the systems described herein may be located on a computer like computer **600**. Similarly, computer executable methods described herein may be performed on a computer like computer **600**. It is to be appreciated that other computers may also be employed with the systems and methods described herein.

[0060] The processor **602** can be a variety of various processors including dual microprocessor and other multi-processor architectures. The memory **604** can include

volatile memory and/or non-volatile memory. The non-volatile memory can include, but is not limited to, read only memory (ROM), programmable read only memory (PROM), electrically programmable read only memory (EPROM), electrically erasable programmable read only memory (EEPROM), and the like. Volatile memory can include, for example,

5 random access memory (RAM), synchronous RAM (SRAM), dynamic RAM (DRAM), synchronous DRAM (SDRAM), double data rate SDRAM (DDR SDRAM), and direct RAM bus RAM (DRRAM).

[0061] A disk 606 may be operably connected to the computer 600 via, for example, an input/output interface (e.g., card, device) 618 and an input/output port 610. The disk 606 can include, but is not limited to, devices like a magnetic disk drive, a solid state disk drive, a

10 floppy disk drive, a tape drive, a Zip drive, a flash memory card, and/or a memory stick. Furthermore, the disk 606 can include optical drives like, a compact disc ROM (CD-ROM), a CD recordable drive (CD-R drive), a CD rewriteable drive (CD-RW drive) and/or a digital video ROM drive (DVD ROM). The memory 604 can store processes 614 and/or data 616, for example. The disk 606 and/or memory 604 can store an operating system that controls 15 and allocates resources of the computer 600.

[0062] The bus 608 can be a single internal bus interconnect architecture and/or other bus or mesh architectures. The bus 608 can be of a variety of types including, but not limited to, a memory bus or memory controller, a peripheral bus or external bus, a crossbar switch,

20 and/or a local bus. The local bus can be of varieties including, but not limited to, an industrial standard architecture (ISA) bus, a microchannel architecture (MSA) bus, an extended ISA (EISA) bus, a peripheral component interconnect (PCI) bus, a universal serial (USB) bus, and a small computer systems interface (SCSI) bus.

[0063] The computer 600 may interact with input/output devices via i/o interfaces 618

25 and input/output ports 610. Input/output devices can include, but are not limited to, a keyboard, a microphone, a pointing and selection device, cameras, video cards, displays, disk 606, network devices 620, and the like. The input/output ports 610 can include but are not limited to, serial ports, parallel ports, and USB ports.

[0064] The computer 600 can operate in a network environment and thus may be

30 connected to network devices 620 via the i/o interfaces 618 and the i/o ports 610. Through the network devices 620, the computer 600 may interact with a network. Through the network, the computer 600 may be logically connected to remote computers. The networks

with which the computer **600** may interact include, but are not limited to, a local area network (LAN), a wide area network (WAN), and other networks. The network devices **620** can connect to LAN technologies including, but not limited to, fiber distributed data interface (FDDI), copper distributed data interface (CDDI), Ethernet/IEEE 802.3, token ring/IEEE 5 802.5, wireless/IEEE 802.11, Bluetooth, and the like. Similarly, the network devices **620** can connect to WAN technologies including, but not limited to, point to point links, circuit switching networks like integrated services digital networks (ISDN), packet switching networks, and digital subscriber lines (DSL).

[0065] Figure 7 illustrates an example image forming device **700** on which the example systems and methods described herein may be implemented. The image forming device **700** may include a memory **710** configured to store print data, for example, or to be used more generally for image processing. The image forming device **700** may include a memory quality assurance logic **715** configured to participate in analyzing the quality of memory locations in memory **710**, in logically removing memory locations known to have errors, and 15 in logically replacing the removed memory locations with replacement memory locations.

[0066] The image forming device **700** may receive print data to be rendered. Thus, the image forming device **700** may include a rendering logic **725** configured to generate a printer-ready image from print data. Rendering varies based on the format of the data involved and the type of imaging device. In general, the rendering logic **725** converts high-level data into a graphical image for display or printing (e.g., the print-ready image). For 20 example, one form is ray-tracing that takes a mathematical model of a three-dimensional object or scene and converts it into a bitmap image. Another example is the process of converting HTML into an image for display/printing. It is to be appreciated that the image forming device **700** may receive printer-ready data that does not need to be rendered and thus 25 the rendering logic **725** may not appear in some image forming devices.

[0067] The image forming device **700** may also include an image forming mechanism **730** configured to generate an image onto print media from the print-ready image. The image forming mechanism **730** may vary based on the type of the imaging device **700** and may include a laser imaging mechanism, other toner-based imaging mechanisms, an ink jet 30 mechanism, digital imaging mechanism, or other imaging reproduction engine. A processor **735** may be included that is implemented with logic to control the operation of the image-forming device **700**. In one example, the processor **735** includes logic that is capable of

executing Java instructions. Other components of the image forming device **700** are not described herein but may include media handling and storage mechanisms, sensors, controllers, and other components involved in the imaging process.

[0068] **Figure 8** illustrates an example operating system transparent system **800** for on-the-fly memory testing. The system **800** includes a memory location identifying logic **810** configured to identify a target memory location and a replacement memory location. Thus, the memory location identifying logic **810** is operably connected to a memory **820**. The memory **820** is accessible via a programmable memory address resolving logic **830** that is configured to provide access to the target memory location and the replacement memory location. The programmable memory address resolving logic **830** can be selectively reprogrammed to divert memory accessing operations from the target memory location to the replacement memory location. Thus, after mirroring the contents of the target memory location to the replacement memory location, memory accessing operations **840** that desire to access the target memory location can be completed by accessing the replacement memory location making the memory testing substantially transparent.

[0069] The system **800** also includes a test controlling logic **850** that is operably connected to the memory location identifying logic **810** and the programmable memory address resolving logic **830**. The test controlling logic **850** may be configured to selectively program the programmable memory address resolving logic **830** to divert memory access operations **840** from the target memory location to the replacement memory location. The test controlling logic **850** may also be configured to initiate memory testing of the target memory location.

[0070] Note that in one example the memory location identifying logic **810**, the programmable memory address resolving logic **830** and the test controlling logic **850** do not consume non-memory operating system resources like processor cycles, process table entries, file table entries, and the like. Thus, the memory testing initiated or performed by the test controlling logic **850** can occur on-the-fly (e.g., while normal system operation is occurring) and will be transparent to (e.g., not halting) memory accessing operations **840** that attempt to access the memory being tested.

[0071] **Figure 9** illustrates an example memory quality assurance method **900**. The method **900** may include, at **910**, selectively copying the contents of a first memory location to a second memory location. For example, the first memory location may be in main

memory and the second memory location may be in main memory, a set-aside buffer, a cache, a temporary memory, and so on. The method **900** may also include, at **920**, logically replacing the first memory location with the second memory location. Thus, memory access requests initially destined for the first memory location will be directed to the second memory
5 location. Thus, the first memory location is logically isolated and available for non-intrusive testing. Thus, the method **900** may also include, at **930**, initiating memory testing of the first memory location without an operating system interaction.

[0072] **Figure 10** illustrates another example memory quality assurance method **1000**. The method **1000** may include, at **1010**, selecting a first memory location to test from a first
10 set of memory. The first set of memory may be, for example, in a main memory. The first memory location may be, for example, an individually addressable unit of memory like a byte, a block, a page, and the like. The method **1000** may also include, at **1020**, selectively copying the contents of the first memory location to a second memory location and, at **1030**, logically replacing the first memory location with the second memory location. This leaves
15 the first memory location logically isolated and thus available for non-intrusive testing (e.g., testing that will be transparent to an operating system that allocates and accesses memory). Thus, the method **1000** may also include, at **1040**, initiating testing of the first memory location.

[0073] **Figure 11** illustrates an example operating system transparent method **1100** for
20 on-the-fly memory testing. The method **1100** may include, at **1110**, identifying a test memory location and a mirroring memory location. The test memory location may be, for example, in a main memory and may be an individually addressable unit of memory like a byte, a block, a page, and the like. The mirroring memory location may be in main memory, a set-aside buffer, a cache, a temporary memory, and so on. The method **1100** may also
25 include, at **1120**, mirroring the test memory location to the mirroring memory location. Thus, the contents of the test memory location are preserved in the mirroring memory location.

[0074] The method **1100** may also include, at **1130**, selectively reconfiguring memory accessing operations so that memory accesses originating in an operating system instance that are addressed to the test memory location are redirected to the mirroring memory location.
30 This leaves the testing memory location logically isolated and available for non-intrusive testing. Thus, the method **1100** may also include, at **1140**, testing the test memory location without disrupting an operating system instance.

[0075] While the systems, methods, and so on have been illustrated by describing examples, and while the examples have been described in considerable detail, it is not the intention of the applicants to restrict or in any way limit the scope of the appended claims to such detail. It is, of course, not possible to describe every conceivable combination of

5 components or methodologies for purposes of describing the systems, methods, and so on employed in memory error ranking. Additional advantages and modifications will readily appear to those skilled in the art. Therefore, the invention, in its broader aspects, is not limited to the specific details, the representative apparatus, and illustrative examples shown and described. Accordingly, departures may be made from such details without departing
10 from the spirit or scope of the applicants' general inventive concept. Thus, this application is intended to embrace alterations, modifications, and variations that fall within the scope of the appended claims. Furthermore, the preceding description is not meant to limit the scope of the invention. Rather, the scope of the invention is to be determined by the appended claims and their equivalents.

15 [0076] To the extent that the term "includes" or "including" is employed in the detailed description or the claims, it is intended to be inclusive in a manner similar to the term "comprising" as that term is interpreted when employed as a transitional word in a claim. Furthermore, to the extent that the term "or" is employed in the claims (e.g., A or B) it is intended to mean "A or B or both". When the applicants intend to indicate "only A or B but
20 not both" then the term "only A or B but not both" will be employed. Thus, use of the term "or" herein is the inclusive, and not the exclusive use. See, Bryan A. Garner, A Dictionary of Modern Legal Usage 624 (2d. Ed. 1995).